Parallelization in LIS

Sujay Kumar
Jim Geiger

Land Information System (LIS)
Annual Team Meeting
21 January 2003

Outline

1. Background
2. Baselining Performance Results
3. Current Design
4. Current Performance Results
5. Planned Work

Background:
Components of LIS

- High resolution LDAS driving multiple LSMs
- Data management system
- Web-based user interface

Background: Main Goal

- A Global Land Data Assimilation System at high resolution (up to 1km) for near real-time land surface modeling

Problem

- Computational requirements significantly increase at domains of high resolution
  - Approximately 2 months to simulate 1 day at 15 minute timestep globally at 1km

Background: Approach/Goals

- Use of scalable, massively parallel computing infrastructure for performance improvement
- Exploit weak horizontal coupling of global land surface processes
- Performance target
  - Achieve 1 ms / gridcell / day at 5 km (Mar 2003)
  - Achieve 0.4 ms / gridcell / day at 1 km (Feb 2004)

Baselining Performance Results

- GLDAS baseline simulations on ESS Testbed were demonstrated as part of July 2002 performance milestone.
  - CLM, NOAH at 1/4 deg.

Baselining Performance Results

- Overall performance
  - CLM: 5–6 ms / gridcell / day
  - NOAH: 5–6 ms / gridcell / day
- Identified "time-hog" routines
  - Time / space interpolation of forcing data
  - Land surface model runs
  - Output

Current Design

- LIS uses a "pool of tasks" approach to exploit loop-based parallelism.
  - Pros
    - Provides automatic load balancing
    - Centralized preprocessing and initializations
    - Platform independent
  - Cons
    - Communication overhead
    - Potential bottleneck at the master node

Current Performance Results

- Parallelized:
  - Time and space interpolation
  - Tile loop of NOAH runs
- Overall performance:
  - Linear speedup in individual routines that are parallelized
  - However, the individual computational tasks are fine grained and hence communication overhead becomes significant enough to offset these improvements
  - Parallel runs $\sim$ 4 ms / gridcell / day at 1/4 degree resolution

Planned Work

- Investigating approaches to reduce communication overheads
  - Use collective communication
  - Move towards a static task distribution model
  - Locally host data on compute nodes